

3

<u>Inventor</u>	<u>Citizenship</u>	<u>Residence Address</u>
Stuart Phillips	United States	
Anthony Li	United States	
David Katz	United States	

10

11

12 The assignee is Cisco Technology, Inc., a California
13 corporation having an office at 170 West Tasman Drive, San Jose,
14 CA 95134.

15

16

Title of the Invention

17

Enhanced Network Services Using a Subnetwork of Communicating
Processors

20

21

Cross-Reference to Related Applications

22

23 This application is a continuation of Application No.
24 09/738,435, filed December 14, 2000, which is a continuation of
25 Application No. 08/582,073, filed January 2, 1996, now Patent No.
26 6,182,224, which claims the benefit of U.S. Provisional Applica-
27 tion No. 60/004,568, filed September 29, 1995, wherein the dis-
28 closure of each application is hereby incorporated by reference.

29

Background of the Invention

1. *Field of the Invention*

This invention relates to enhanced network services using a subnetwork of communicating processors.

2. *Description of Related Art*

A computer network having more than one possible message delivery path, such as for example a network of computer networks (an "internetwork"), generally comprises a set of processors which have the task of routing each message from its source to its destination. These processors (herein called "routers", but sometimes called "bridges", "brouters", or other terms) generally interoperate using a common distributed routing technique. For example, (1) routers generally exchange information about routing conditions in the network, so as to globally inform all routers of conditions which are detected locally, (2) routers generally each route messages according to the same techniques using the same routing technique.

The collection of routers (herein called the "subnet" of routers) thus has access to a rich collection of distributed information about the status of the network needed for routing messages therein, including detailed and up-to-date information about the network topology, relative distance measures in the network, administrative policies in force in the network (such

1 as, for example, network "firewalls" or routers which will route
2 only a subset of messages), and other information. It would be
3 advantageous to provide at least some of this distributed infor-
4 mation to applications which might use it to provide enhanced
5 services. Examples of such enhanced services, and how they would
6 be provided, are described herein.

7

8 The router subnet also provides a powerful and avail-
9 able resource for distributing other information, not strictly
10 required for the task of routing messages in the network, such
11 as, for example, relative load at each host in the network. It
12 would be advantageous to disseminate this information and provide
13 at least some of it to applications which might use it to provide
14 enhanced services. Examples of such enhanced services, and how
15 they would be provided, are described herein.

16

17 Accordingly, it would be advantageous to provide a
18 method and system for enhanced network services using a subnet-
19 work of communicating processors.

20

21 The following U.S. Patent(s) or other documents may be
22 pertinent:

23

24 o R. Srinivasan, "RPC: Remote Procedure Call Protocol
25 Specification Version 2", Network Working Group RFC
26 1831 (August 1995).

27

28 The pertinence of the related art will also be apparent
29 to those skilled in the art after perusal of this application.

Summary of th Invention

The invention provides a method and system for providing enhanced services for a network, using a subnetwork of communicating processors. The enhanced services use information about the network which is available to the subnet of communicating processors (such as a set of routers), interoperating using a common distributed technique for disseminating that network information. In a first aspect of the invention, the router subnet collects network topology information and provides a service using that network topology information, responsive to requests from non-routers (such as host processors) coupled to the network. The network topology information comprises information about paths and routes, including bandwidth, connectivity, delay, traffic reservations, and administrative policies applicable to those paths and routes. Routers providing the enhanced service have the option of requiring authentication for service requests.

For a first example, the router subnet provides an enhanced distributed naming service which, in addition to translating server names into host addresses, also orders those host addresses by relative distance in the network, or by another criterion designated by the client. For a second example, the router subnet provides an enhanced message delivery service which, in addition to delivering a message to a plurality of destinations, assures that all destinations receive the message at substantially the same time.

1

2 In a second aspect of the invention, the router subnet
3 collects information advertised by hosts coupled to the network,
4 and disseminates that host information to substantially all
5 routers, using the common distributed technique for disseminating
6 network topology information. In a first preferred embodiment,
7 the host information comprises information about server processes
8 available at the originating host (such as what services are
9 available and to which users those services are available). In a
10 second preferred embodiment, the host information comprises in-
11 formation about client processes operating at the originating
12 host (such as which users are operating those client processes
13 and which services they desire).

14

15 For a first example, hosts advertise their relative
16 load for dissemination by the router subnet, and the router sub-
17 net provides an enhanced distributed naming service which, in ad-
18 dition to translating server names into host addresses, also or-
19 ders those host addresses by relative load (or by both relative
20 distance and relative load). Similarly, hosts may also advertise
21 other server information, such as cost for performing the ser-
22 vice, delay in performing the service, or other administrative
23 policies which would affect the choice of server.

24

25 For a second example, hosts advertise their logged-in
26 users for dissemination by the router subnet, and the router sub-
27 net provides an enhanced message-delivery service which, in addi-
28 tion to receiving messages for delivery, also delivers those mes-
29 sages to the host where the destination user is logged-in. Simi-

larly, hosts may also advertise other client information, such as willingness to pay for performing a designated service, which the router subnet may disseminate for responses by servers or by other clients.

Brief Description of the Drawings

Figure 1A shows a block diagram of a system for providing enhanced services, and also shows a block diagram of a system for providing enhanced services using server advertisements, in a computer network. Figure 1B shows a flow diagram of a method for providing enhanced services, and also shows a flow diagram of a method for providing enhanced services using server advertisements in a computer network as shown in figure 1A.

Figure 2A shows a block diagram of a system for providing synchronized message receipt in a computer network. Figure 2B shows a flow diagram of a method for providing synchronized message receipt in a computer network as shown in figure 2A.

Figure 3A shows a block diagram of a system for providing enhanced services using both client and server advertisements, in a computer network. Figure 3B shows a flow diagram of a method for providing enhanced services using both client and server advertisements, in a computer network as shown in figure 3A.

1 Figure 4A shows formats for a set of packets for use
2 with an authenticated remote procedure call transport ("ART")
3 protocol. Figure 4B shows a flow diagram for an authenticated
4 remote procedure call transport protocol.

6 **Description of the Preferred Embodiment**

8 In the following description, a preferred embodiment of
9 the invention is described with regard to preferred process steps
10 and data structures. However, those skilled in the art would
11 recognize, after perusal of this application, that embodiments of
12 the invention may be implemented using a set of general purpose
13 computers operating under program control, and that modification
14 of a set of general purpose computers to implement the process
15 steps and data structures described herein would not require un-
16 due invention.

18 **PROVIDING ENHANCED NETWORK SERVICES**

20 Figure 1A shows a block diagram of a system for provid-
21 ing enhanced services, in a computer network. Figure 1B shows a
22 flow diagram of a method for providing enhanced services, in a
23 computer network as shown in figure 1A.

25 In an internetwork 100 to which a network 101 is cou-
26 pled, a client process 102 on a client node 103 contacts a server
27 process 104 on a server node 105, to obtain data or a service.
28 For example, the data may comprise a file located at the server
29

1 node 105; the service may comprise conducting a computation, pro-
2 viding data in response to search parameters, or even just pro-
3 viding the current time.

4
5 The verb "contacts" refers to communication between the
6 client process 102 and the server process 104 using messages com-
7 prising message packets which are transmitted using the internet-
8 work 100 between the client process 102 and the server process
9 104. Messages and message packets are transmitted using the in-
10 ternetnetwork 100 using a subnet 110 of routers 111. Each router
11 111 executes a routing technique which causes the routers 111 to
12 collectively interoperate using a common distributed technique
13 for routing messages and distributing information about the in-
14 ternetnetwork 100. This network information is stored at each
15 router 111 in one or more network tables 112, and generally com-
16 prises information about paths and routes, such as bandwidth,
17 connectivity, delay, traffic reservations, and administrative
18 policies applicable to those paths and routes.

19
20 Particular routing techniques, such as IGRP, AppleTalk,
21 and IPX, are known in the art of networking.

22
23 A method 150 is conducted in cooperation by various
24 nodes and processes coupled to the internetwork 100. At a flow
25 point 160 for the method 150, the client process 102 desires to
26 contact a server process 104, the latter of which is coupled to
27 the internetwork 100.

1 At a step 161, the client process 102 first contacts a
2 name server 105, which preferably comprises a server process 104
3 executing on a server node 103 for which the client process 102
4 already has a network address, using a name-request message 121.
5 In this example, the name server 105 is shown executing on a dif-
6 ferent node and coupled to the same network 101 as the client
7 node 103, but there is no particular requirement that the name
8 server 105 is in fact so disposed. For example, the name server
9 105 may be located at the same node as the client process 102, or
10 may be coupled to a different network 101 on the internetwork
11 100, so long as the client process 102 has available to it a suf-
12 ficient network address to transmit the name-request message 121
13 to the name server 105.

14

15 In a preferred embodiment, the name server 105 com-
16 prises a DNS name server operating according to the DNS protocol.

17

18 The name-request message 121 asks the name server 105
19 to translate an identifier for the service provided into a net-
20 work address for a server process 104 providing that service.
21 For example, the identifier may comprise a host name for a node
22 coupled to the internetwork 100, such as "sanjose.cisco.com", and
23 the name server 105 may translate that host name into an IP ad-
24 dress for that node, such as "3.141.59.26".

25

26 At a step 162, the name server 105 receives the name-
27 request message 121, parses the name-request message 121 to de-
28 termine the name to be translated, looks up the name in a table
29 it maintains for translating names into network addresses, and

1 transmits a name-response message 122 comprising one or more net-
2 work addresses to the client process 102 at the client node 103.

3
4 In general, the name to be translated may specify a
5 service which is provided by more than one server process 104 or
6 by more than one server node 105, so that the name server 105 has
7 more than one network address to provide in the name-response
8 message 122. In a preferred embodiment, the name server 105 pro-
9 vides all such network addresses in the name-response message
10 122. Order is not critical at this flow point of the method 150.

11
12 At a step 163, the client process 102 receives the
13 name-response message 122 and parses the name-response message
14 122 to determine the network addresses provided by the name
15 server 105. The client process 102 then contacts a router 111
16 using an order-request message 123.

17
18 The order-request message 123 asks the router 111 to
19 order a set of addresses by their relative distance on the inter-
20 network 100. Preferably, the client process 102 packages the ad-
21 dresses from the name-response message 122 into the order-request
22 message 123.

23
24 At a step 164, the router 111 receives the order-
25 request message 123, parses the order-request message 123 to de-
26 termine the set of addresses, looks up each address in its net-
27 work tables 112 to determine a relative distance for that network
28 address, orders the network addresses by relative distance, and

1 transmits an order-response message 124 comprising the network
2 addresses ordered relative distance.

3

4 The relative distance for each network address may dif-
5 fer at different times, so the order provided by the router 111
6 in the order-response message 124 may also differ in response to
7 order-request messages 123 transmitted at different times.

8

9 In a preferred embodiment, the relative distance for
10 each network address is computed in response to the expected time
11 delay in delivering a message from the client node 103 to the
12 server node 105. It is expected that this expected time delay
13 will be nearly the same as the expected time delay from the name
14 server 105. In alternative embodiments described below in which
15 the name server 105 orders the addresses in response to relative
16 distance information from the router 111, the ordering may turn
17 out to differ from an ordering for the client node 103, but this
18 is not expected to be important unless the name server 105 is at
19 a substantial relative distance from the client node 103.

20

21 In alternative embodiments, the relative distance for
22 each network address may also be computed in response to other
23 information available to the router 111 in its network tables
24 112, including for example (1) an expected bandwidth of a path
25 from the client node 103 to the server node 105, (2) any traffic
26 reservations on paths between the client node 103 to the server
27 node 105, or (3) any administrative policies applicable to those
28 paths between the client node 103 to the server node 105. For
29 example, if a required path to one of the server nodes 105 is

1 known to the router 111 to be heavily trafficked between the
2 hours of 9:00 a.m. and 5:00 p.m., that server node 105 may be as-
3 signed a substantially higher relative distance responsive to
4 that information.

5
6 Although in a preferred embodiment, the name-request
7 message 121 and the order-request message 123 are both preferably
8 transmitted by the client process 102, in alternative embodi-
9 ments, alternative arrangements are possible which would obtain
10 the same information for the client process 102. Such alterna-
11 tive arrangements would be clear to those skilled in the art af-
12 ter perusing this application, and are within the scope and
13 spirit of the invention. In a first alternative embodiment, the
14 name server 105 may obtain an ordering of network addresses from
15 the router 111, and may order the network addresses in the name-
16 response message 122. To do so, the name server 105 may transmit
17 an order-request message 123 to the router 111 and use the re-
18 sults it receives in the order-response message 124. In a second
19 alternative embodiment, the name server 105 may be located at the
20 router 111, and may order the network addresses in the name-
21 response message 122 in response to direct access to the network
22 tables 112.

23
24 In a preferred embodiment, the network address for one
25 or more particular server nodes 105 may be restricted to those
26 client processes 102 with sufficient authorization. In such a
27 case, the name server 105 would require that the client process
28 102 is authenticated, to determine that the name-request message
29 121 really comes from the client process 102 and not someone

1 else, and that the client process 102 has sufficient permissions,
2 in particular, that the name server 105 has been informed that
3 the client process 102 is authorized to receive that information.
4 A preferred authentication process is described below with refer-
5 ence to figure 4.

6
7 In a preferred embodiment, the relative distance or
8 other ordering information may also be restricted to those client
9 processes 102 with sufficient authorization. In such a case, the
10 router 111 would require that the client process 102 is authenti-
11 cated, to determine that the name-request message 121 really
12 comes from the client process 102 and not someone else, and that
13 the client process 102 has sufficient permissions, in particular,
14 that the router 111 has been informed that the client process 102
15 is authorized to receive that information.

16
17 In some cases, the router 111 may have insufficient in-
18 formation to respond to the order-request message 123 on its own.
19 In such a case, the router 111 may forward the order-request mes-
20 sage 123 (or another message with similar information) to a sec-
21 ond router 111. The second router 111 may also find that it must
22 forward the order-request message 123 or other message, until one
23 or more routers are reached which, individually or collectively,
24 do have sufficient information to respond to the original order-
25 request message 123.

26
27 At a step 165, the client process 102 selects a first
28 one of the network addresses, preferably the network address with
29 the smallest relative distance, and contacts the server process

1 104 on the server node 105 with that network address. In the
2 case that the client process 102 is unable to contact the server
3 process 104 at the first network address, the client process 102
4 selects another, preferably the network address with the next
5 smallest relative distance, and contacts the server process 104
6 on the server node 105 with that network address. The client
7 process 102 repeats this step until one of the server processes
8 104 is contacted or the set of network addresses is exhausted.

AUTOMATED RESOURCE DISTRIBUTION

12 A first example enhanced network service is automated
13 distribution of resources such as services. Some services are
14 readily distributed by providing more than one service provider,
15 each of which can fully provide substantially the same service.
16 For example, a service providing a collection of data may be du-
17 plicated by copying the data to multiple sites; a service provid-
18 ing a program to be executed may be duplicated by copying the
19 program to multiple sites, and providing resources for executing
20 the program at each of those sites.

22 In this example enhanced service, a resource to be
23 automatically distributed is embodied in server processes 104 at
24 a plurality of server nodes 105. For example, the resource may
25 comprise one of the following:

1 o an FTP directory, such as a directory of files which may be
2 copied from the server node 105 to the client node 103 using
3 the "FTP" protocol;

4
5 o a network news server, such as a server for providing mes-
6 sages posted to "network news" or to a bulletin board sys-
7 tem, and distributed substantially throughout a portion of
8 the internetwork 100;

9
10 o a World Wide Web page, such as a file comprising a set of
11 data to be displayed, possibly including text, graphics,
12 display commands, or motion picture data, and comprising one
13 or more hypertext links to other such files;

14
15 o a database and search engine, such as a program for receiv-
16 ing search requests, executing those search requests against
17 a database, and providing the results of such search re-
18 quests; or

19
20 o a time server, such as a program for providing a time of day
21 or similar globally available information, such as a dic-
22 tionary server, spell-checker or thesaurus.

23
24 Other and further possibilities will readily occur to
25 those skilled in the art after perusal of this application; such
26 other and further possibilities would not require invention or
27 undue experiment, and are within the scope and spirit of the in-
28 vention.

1 The name servers 105 which translate the service name
2 into network addresses are given the network addresses for the
3 plurality of server nodes 105 at which the server processes 104
4 may be found. This information is distributed to the name serv-
5 ers 105 using a name server protocol such as the DNS protocol.
6 Name server protocols are known in the art of networking, and so
7 are not discussed in detail herein.

8
9 Duplicating the server process 104 and distributing
10 multiple network addresses for the service to the name servers
11 105 has the following effects: When the name-request message 121
12 is transmitted at the step 161 from the client process 102 to the
13 name server 105, the name server 105 obtains a plurality of net-
14 work addresses for server processes 104 at a plurality of server
15 nodes 105. When the order-request message 123 is transmitted at
16 the step 163 from the client process 102 to the router 111, the
17 router 111 orders the plurality of network addresses for server
18 processes 104 at a plurality of server nodes 105 by their rela-
19 tive distance from the client node 103. When the client process
20 102 contacts a server process 104 using a network address from
21 the order-response message 124, the server process 104 which is
22 chosen is the one at the server node 105 with the smallest rela-
23 tive distance from the client node 103.

24
25 Thus, when the client process 102 initiates the method
26 150 to contact one of the server processes 104, it will automati-
27 cally be assigned to the relatively nearest server process 104;
28 when multiple client process 102 attempt to contact one of the
29

server processes 104, they will automatically be distributed among the set of server processes 104 which may be chosen.

SYNCHRONIZED MESSAGE RECEIPT

Figure 2A shows a block diagram of a system for providing synchronized message receipt in a computer network. Figure 2B shows a flow diagram of a method for providing synchronized message receipt in a computer network as shown in figure 2A.

A second example enhanced network service is synchronized distribution of messages to recipients, such as (1) assuring that a designated message is delivered to a recipient node on the internetwork 100 at a designated time, or (2) assuring that a designated message is delivered to multiple recipients at the same time.

In this example enhanced service, a message is labeled with a set of designated destination nodes and a designated time for delivery to those destination nodes. For example, the message may comprise one of the following:

- o a financial news report, such as a ticker-tape entry which is intended for delivery to customers fifteen minutes after its occurrence in real time; or

o a wake-up call or other reminder, such as a greeting message which is intended for delivery at a time specified by the recipient.

In a preferred embodiment, each one of the routers 111 maintains an indicator of the current time, synchronized across all routers 111 so that any pair of routers 111 will have the same value for the "current time", preferably within an error margin of less than about 0.5 milliseconds. This indicator may be recorded in the network tables 112 or in another register within each router 111. The routers 111 may maintain this indicator by means of a network time protocol, or by some other technique. Network time protocols, which are preferred, are known in the art of networking.

A method 250 is conducted in cooperation by the client process 102 and various routers 111 on the internetwork 100. At a flow point 260 for the method 250, the client process 102 desires to send a message 200 for receipt at a destination node 201 at a designated time.

At a step 261, the client process 102 contacts one of the routers 111 using a timed message 210. Since all messages on the internetwork 100, other than those limited to the local network 101, are transmitted using one or more routers 111, the client process 102 need only transmit the timed message 210 to the destination node 201 as if it were to transmit any other message to the destination node 201, and the timed message 210 will be received by one of the routers 111 along a path to the destina-

tion node 201. The timed message 210 comprises an ordinary message for delivery to, plus a header designating the desired time for receipt at, the destination node 201.

At a step 262, one of the routers 111 along a path to the destination node 201 receives the timed message 210 and parses it to determine the destination node 201, and determines if it would be one of the routers 111 to effect final delivery of the message to the destination node 201. If there is a single destination node 201, there will be only one such router 111; if there are multiple destination nodes 201, there may be more than one such router 111. If this router 111 would not effect final delivery, the method 250 continues with the step 263; if this router 111 would effect final delivery, the method 250 continues with the step 264.

At a step 263 (the router 111 would not effect final delivery), the router 111 forwards the timed message 210 to the next router 111 in the path to the destination node 201, just like any other message. In a preferred embodiment, if the desired time for receipt is very soon, the router 111 may reply to the client process 102 with an error message, indicating that delivery at the designated time is impossible, unlikely, or merely not guaranteed. In alternative embodiments, if the desired time for receipt is very soon, the router 111 may give priority to the timed message 210 over other messages.

At a step 264 (the router 111 would effect final delivery), the router 111 holds the timed message 210 until the design-

1 nated time for delivery to the destination node 201, and at that
2 time, delivers the message just like any other message. In a
3 preferred embodiment, if the current time is past the desired
4 time for receipt, the router 111 may reply to the client process
5 102 with an error message, indicating that delivery at the desig-
6 nated time did not occur, and may either discard the message or
7 deliver it late. In alternative embodiments, if the desired time
8 for receipt is very soon, the router 111 may give priority to the
9 timed message 210 over other messages.

10
11 In a preferred embodiment, the capability to send a
12 timed message 210 may be restricted to those client processes 102
13 with sufficient authorization. In such a case, the first router
14 111 to receive and parse a timed message 210 would require that
15 the client process 102 is authenticated, to determine that the
16 timed message 210 really comes from the client process 102 and
17 not someone else, and that the client process 102 has sufficient
18 permissions, i.e., that the router 111 has been informed that the
19 client process 102 is authorized to send timed messages 210.

20 21 **ENHANCED SERVICES USING SERVER ADVERTISEMENTS**

22
23 Figure 1A also shows a block diagram of a system for
24 providing enhanced services using server advertisements, in a
25 computer network. Figure 1B also shows a flow diagram of a
26 method for providing enhanced services using server advertise-
27 ments, in a computer network as shown in figure 1A.

1 In a preferred embodiment, the method 150 is supple-
2 mented by using information "advertised" by the server processes
3 104. At a flow point 170 for the supplemented method 150, one of
4 the server processes 104 desires to disseminate server advertis-
5 ing information to the subnet 110 of routers 111 for use with an
6 enhanced network service.

7
8 The server advertising information comprises any infor-
9 mation provided by one of the server processes 104 for dissemina-
10 tion to the subnet 110 of routers 111. For example, server ad-
11 vertising information may comprise the following:

- 12
- 13 o relative load at the server node;
 - 14
 - 15 o resources available at the server node, such as what ser-
16 vices are available and to which classes of client process
17 102 those services are available, amount of free mass stor-
18 age space or free processor time;
 - 19
 - 20 o services available at the server node, to which classes of
21 users those services are available, cost for performing
22 those services, expected delay in performing those services,
23 or other administrative policies which could affect the
24 choice of server process 104 by a client process 102; or
 - 25
 - 26 o which persons are logged in at the server node 105.
- 27

28 At a step 171, one of the server processes 104 sends a
29 server-info message 131 to one of the routers 111.

1
2 The routing technique is supplemented by providing that
3 each one of the routers 111 disseminates information from the
4 server-info message 131 to all routers 111 in the subnet 110.
5 Each router 111 which receives the server-info message 131 per-
6 forms the steps 172, 173, and 174. As a result, each router 111
7 eventually receives the information to be disseminated in the
8 server-info message 131 and stores that information in its net-
9 work tables 112.

10

11 At a step 172, the router 111 receives the server-info
12 message 131 and parses the server-info message 131 to determine
13 the information to be disseminated.

14

15 At a step 173, the router 111 stores the information to
16 be disseminated in its network tables 112.

17

18 At a step 174, the router 111 forwards the information
19 to be disseminated to a set of neighboring routers 111 in the
20 subnet 110.

21

22 In the supplemented method 150, the step 164 is supple-
23 mented to determine the order for the network addresses respon-
24 sive to the server advertising information. Thus for example, at
25 the step 164 the router 111 orders the network addresses respon-
26 sive to relative load on the server node 105, or responsive to
27 both relative distance to the server node 105 and relative load
28 on the server node 105.

29

1 In other enhanced network services, the routers 111 may
2 route messages to server processes 104 responsive to the server
3 advertising information. For a first example, the server adver-
4 tising information may comprise a list of persons who are cur-
5 rently logged in at the server node 105, and the routers 111 may
6 receive and route a person-to-person message naming a designated
7 person to the server node 105 at which that designated person is
8 logged in. For a second example, the server advertising informa-
9 tion may comprise a set of administrative policies regarding
10 traffic the server node 105 desires not to retransmit, and the
11 routers 111 may receive and route such messages away from such
12 server nodes 105.

14 **ENHANCED SERVICES USING BOTH CLIENT AND SERVER ADVERTISEMENTS**

16 Figure 3A shows a block diagram of a system for provid-
17 ing enhanced services using both client and server advertise-
18 ments, in a computer network. Figure 3B shows a flow diagram of
19 a method for providing enhanced services using both client and
20 server advertisements, in a computer network as shown in figure
21 3A.

23 A method 350 is conducted in cooperation by various
24 client processes 102, client nodes 103, server processes 104, and
25 server nodes 105 and coupled to the internetwork 100. At a flow
26 point 360 for the method 350, the client process 102 desires to
27 disseminate client advertising information to the subnet 110 of
28 routers 111 for use with an enhanced network service.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29

The client advertising information comprises any information provided by one of the client processes 102 for dissemination to the subnet 110 of routers 111. For example, client advertising information may comprise the following:

- o relative urgency of requests from the client process 102; or
- o willingness to pay for services from the server node 105.

At a step 361, one of the client processes 102 sends a client-info message 301 to one of the routers 111, similar to the step 171.

The routing technique is supplemented by providing that each one of the routers 111 disseminates information from the client-info message 301 to all routers 111 in the subnet 110. Each router 111 which receives the client-info message 301 performs the steps 362, 363, and 364. As a result, each router 111 eventually receives the information to be disseminated in the client-info message 301 and stores that information in its network tables 112.

At a step 362, the router 111 receives the client-info message 301 and parses the client-info message 301 to determine the information to be disseminated, similar to the step 172.

1 At a step 363, the router 111 stores the information to
2 be disseminated in its network tables 112, similar to the step
3 173.

4
5 At a step 364, the router 111 forwards the information
6 to be disseminated to a set of neighboring routers 111 in the
7 subnet 110, similar to the step 174.

8
9 At a flow point 370 for the method 350, the server
10 process 102 desires to disseminate server advertising information
11 to the subnet 110 of routers 111 for use with an enhanced network
12 service.

13
14 The routing technique is also supplemented by providing
15 that each one of the routers 111 disseminates information from
16 the server-info message 131 to all routers 111 in the subnet 110,
17 similar to the supplemented method 150. Each router 111 which
18 receives the server-info message 131 performs the steps 372, 373,
19 and 374. As a result, each router 111 eventually receives the
20 information to be disseminated in the server-info message 131 and
21 stores that information in its network tables 112.

22
23 At a step 372, the router 111 receives the server-info
24 message 131 and parses the server-info message 131 to determine
25 the information to be disseminated, similar to the step 172.

26
27 At a step 373, the router 111 stores the information to
28 be disseminated in its network tables 112, similar to the step
29 173.

1
2 At a step 374, the router 111 forwards the information
3 to be disseminated to a set of neighboring routers 111 in the
4 subnet 110, similar to the step 174.

5
6 At a flow point 380, one of the routers 111 receives
7 both client advertising information and server advertising infor-
8 mation.

9
10 At a step 381, the router 111 matches the client adver-
11 tising information against the server advertising information.

12
13 The nature of this step 381 depends on the nature of
14 the client advertising information and the server advertising in-
15 formation. For example, if the client advertising information
16 indicates the desire to obtain a named service X for a price Y or
17 less, and the server advertising information indicates the abil-
18 ity to provide the named service X for a price Y or more, the
19 router 111 obtains a match. Other situations in which the router
20 111 might obtain a match include the following:

- 21
22 o the client process 102 has an indicated degree of urgency
23 and the server process 104 has an indicated degree of ex-
24 pected delay to perform the named service;
25
26 o the client process 102 has an indicated priority and the
27 server process 104 has an indicated degree of free capacity
28 to perform the named service;

1 o the client process 102 has an indicated set of server re-
2 quirements and the server process 104 has an indicated set
3 of server capabilities to perform the named service;

4
5 o the client process 102 has an indicated set of subtasks of
6 the named service to be performed and the server process 104
7 has an indicated set of subtasks of the named service it is
8 capable of performing; or

9
10 o the client process 102 has an indicated set of products or
11 services a first associated user desires to obtain and the
12 server process 104 has an indicated set of products or ser-
13 vices a second associated user desires to provide, such as
14 for example airline tickets, computer equipment, office sup-
15 plies, stocks or other securities, or other products or ser-
16 vices.

17
18 At a step 382, the router 111 transmits a client-match
19 message 302 to the client process 102, and a server-match message
20 303 to the server process 104, indicating the nature of the
21 match.

22
23 At a step 383, the client process 102 and the server
24 process 104 communicate about the matched information.

AUTHENTICATED REMOTE PROCEDURE CALL TRANSPORT

Figure 4A shows formats for a set of packets for use with an authenticated remote procedure call transport ("ART") protocol. Figure 4B shows a flow diagram for an authenticated remote procedure call transport protocol.

The ART protocol described herein comprises four types of packets: (1) an ART-request packet, indicating a request for a service to be performed at the remote site, (2) an ART-reply packet, indicating a response to the ART-request packet, (3) an ART-auth-challenge packet, indicating a challenge to the recipient to show that the recipient is authorized to make the request, and (4) an ART-auth-response packet, indicating a response to the authorization challenge.

A single transaction comprises the sequence (1) ART-request, (2) ART-reply. If authentication is performed, a single transaction comprises the sequence (1) ART-request, (2) ART-auth-challenge, (3) ART-auth-response, (4) ART-reply.

All four types of packet comprise a fixed length packet header 400 and variable length packet data 410. The packet header 400 comprises a version number 401, a packet type value 402, a reply port value 403, and a transaction identifier 404.

The version number 401 comprises one octet (eight-bit byte) specifying the version of the remote procedure call protocol.

1
2 The packet type value 402 comprises one octet specify-
3 ing the type of operation the packet is for. In a preferred em-
4 bodiment, a request packet has the value 1; a reply packet has
5 the value 2; an authentication challenge packet has the value 3;
6 an authentication response packet has the value 4.

7
8 The reply port value 403 comprises two octets specify-
9 ing a 16-bit value of the port at which the client process 102
10 expects to receive further ART protocol packets. The client sets
11 this value with its first ART-request packet; thereafter the cli-
12 ent and server copy this value to all packets for a single ART
13 transaction.

14
15 The transaction identifier 404 comprises four octets
16 specifying a 32-bit unsigned integer identifier for a single
17 transaction. The client sets this value with its first ART-
18 request packet for each transaction; thereafter the client and
19 server copy this value to all packets for a single ART transac-
20 tion. In a preferred embodiment, the client chooses a random
21 nonzero value for its first ART transaction and increments this
22 value by one for each subsequent ART transaction; zero is a valid
23 value which may occur by overflow.

24
25 For the ART-request packet, the variable length packet
26 data 410 comprises one or more requests. Each request comprises
27 a mode/type field 411, a length value 412, and a data field 413.

1 The mode/type field 411 comprises two octets specifying
2 a one-bit authentication bit and a 15-bit request type.

3
4 The length value 412 comprises two octets specifying a
5 16-bit count of the number of octets in the data field 413.

6
7 The data field 413 comprises the actual data for the
8 request. Its interpretation may vary depending on the type of
9 request.

10
11 For the ART-auth-challenge packet, the variable length
12 packet data 410 comprises one or more challenges, one for each
13 request in the ART-request packet requiring authorization. Each
14 challenge comprises the mode/type field 411, the length value
15 412, and the data field 413, formatted like those for requests in
16 the ART-request packet.

17
18 The server copies the mode/type field 411 from the
19 mode/type field 411 for the ART-request packet.

20
21 The data field 413 comprises a randomly chosen chal-
22 lenge value.

23
24 For the ART-auth-response packet, the variable length
25 packet data 410 comprises one response for each challenge in the
26 ART-auth-challenge packet. Each response comprises the mode/type
27 field 411, the length value 412, and the data field 413, format-
28 ted like those in the ART-request packet.

1 The client copies the mode/type field 411 from the
2 mode/type field 411 for the ART-auth-challenge packet.

3

4 The data field 413 comprises the response value appro-
5 priate to the corresponding challenge value.

6

7 For the ART-reply packet, the variable length packet
8 data 410 comprises one reply for each request in the ART-request
9 packet. Each reply comprises the mode/type field 411, the length
10 value 412, and the data field 413, formatted like those in the
11 ART-request packet, and a result field 414.

12

13 The server copies the mode/type field 411 from the
14 mode/type field 411 for the ART-request packet.

15

16 The result field 414 comprises two octets specifying a
17 16-bit value indicating the success or failure of the correspond-
18 ing request. If the request is successful, the value will be
19 zero. If the request is not successful, the value will be non-
20 zero, and particular nonzero values may (at the server's option)
21 indicate the type of failure.

22

23 The ART protocol 450 is conducted in cooperation by the
24 client process 102 and the server process 104. At a flow point
25 460 for the method 450, the client process 102 desires to initi-
26 ate one or more ART requests.

27

28 At a step 461, the client process 102 constructs an
29 ART-request packet 470 for transmission to the server process

1 104, and transmits the ART-request packet 470 to the server proc-
2 ess 104. The ART-request packet 470 comprises one or more re-
3 quests, as described above. The client process 102 sets a timer
4 to protect against the possibility that the packet will be lost.
5 A preferred timer value is 30 seconds. If no reply is received,
6 the client process 102 retries the request, up to five times.
7 ART requests are therefore preferably idempotent.

8
9 At a step 462, the server process 104 receives the ART-
10 request packet 470 and parses the ART-request packet 470 to de-
11 termine each of the requests. If authentication is required for
12 one or more of the requests, the protocol 450 continues with the
13 step 463. If authentication is not required for any of the re-
14 quests, the protocol 450 continues with the step 465. It is pos-
15 sible that authentication will be required for some but not all
16 of the requests in a single ART-request packet 470.

17
18 At a step 463, the server process 104 constructs an
19 ART-auth-challenge packet 471 for transmission to the client
20 process 102, and transmits the ART-auth-challenge packet 471 to
21 the client process 102. The server process 104 sets a timer to
22 protect against the possibility that the packet will be lost. A
23 preferred timer value is 30 seconds. If no response to the chal-
24 lenge is received, the server process 104 need not retry the
25 challenge.

26
27 At a step 464, the client process 102 receives the ART-
28 auth-challenge packet 471 from the server process 104, parses the
29 ART-auth-challenge packet 471 to determine the challenge data,

1 determines the appropriate authorization response, constructs an
2 ART-auth-response packet 472 for transmission to the server proc-
3 ess 104, and transmits the ART-auth-response packet 472 to the
4 server process 104. The client process 104 sets a timer to pro-
5 tect against the possibility that the packet will be lost. A
6 preferred timer value is 30 seconds. If no reply to the original
7 request is received, the client process 104 retries the original
8 request, up to file times.

9

10 A preferred embodiment uses the shared secret model of
11 authentication, as described in RFC-1344, hereby incorporated by
12 reference as if fully set forth herein. The server process 104
13 generates a random challenge value for each request requiring au-
14 thentication; the client process 102 and server process 104 each
15 compute a hash function applied to the concatenation of the chal-
16 lenge value and the shared secret. A preferred embodiment uses
17 the MD5 hash function, as described in RFC-1321, hereby incorpo-
18 rated by reference as if fully set forth herein. The client
19 process 102 transmits the computed value as the response value to
20 the server process 104, which matches the response value against
21 its own computed value. Authentication succeeds if the response
22 value matches the server's own computed value.

23

24 Those skilled in the art will recognize, after perusal
25 of this application, that other and further types of authentica-
26 tion may be used instead or in addition.

27

28 At a step 465, the server process 104 receives the ART-
29 auth-response packet 472 from the client process 102, parses the

1 ART-auth-response packet 472 to determine the response data, and
2 determines whether the authorization challenge was passed by the
3 response. If the client process 102 passed the authorization
4 challenge, the server process 104 performs the request, and the
5 protocol continues with the step 466. If the client process 102
6 failed to pass the authorization challenge, the server process
7 104 generates an failed-authorization error, and the protocol
8 continues with the step 466.

9
10 At a step 466, the server process 104 constructs an
11 ART-reply packet 473 for transmission to the client process 102,
12 and transmits the ART-reply packet 473 to the client process 102.

13
14 At a step 467, the client process 102 receives the ART-
15 reply packet 473 from the server process 104 and parses the ART-
16 reply packet 473 to determine the reply for each request. The
17 client process 102 matches the transaction identifier for the
18 ART-reply packet 473 to a list of outstanding requests; if there
19 is no match, the client process 102 discards the ART-reply packet
20 473.

21 22 ***Alternative Embodiments***

23
24 Although preferred embodiments are disclosed herein,
25 many variations are possible which remain within the concept,
26 scope, and spirit of the invention, and these variations would
27 become clear to those skilled in the art after perusal of this
28 application.